arXiv:1611.01230v1 [stat.ML] 4 Nov 2016

Manuscript submitted to
AIMS' Journals
Volume **X**, Number **0X**, XX **200X**

doi:10.3934/xx.xx.xx.xx

pp. **X–XX**

STATISTICAL INVERSE FORMULATION OF OPTICAL FLOW WITH UNCERTAINTY QUANTIFICATION

Jie Sun* and Erik Bollt

Department of Mathematics, Clarkson University
Potsdam, NY 13699-5815, USA

(Communicated by the associate editor name)

Abstract. Optical flow refers to the visual motion observed between two consecutive images. Since the degree of freedom is typically much larger than the constraints imposed by the image observations, the straightforward formulation of optical flow inference is an ill-posed problem. By setting some type of additional "regularity" constraints, classical approaches formulate a well-posed optical flow inference problem in the form of a parameterized set of variational equations. In this work we build a mathematical connection, focused on optical flow methods, between classical variational optical flow approaches and Bayesian statistical inversion. A classical optical flow solution is in fact identical to a maximum a posteriori estimator under the assumptions of linear model with additive independent Gaussian noise and a Gaussian prior distribution. Unlike classical approaches, the statistical inversion approach to optical flow estimation not only allows for "point" estimates, but also provides a distribution of solutions which can be used for ensemble estimation and in particular uncertainty quantification.

1. **Introduction.** Optical flow reflects the visual motion between consecutive images. Determination of optical flow is important for applications ranging from machine learning and computer vision [22], artificial intelligence and robotics [4, 21], to scientific applications from oceanography to weather forecasting [6, 5, 8, 16, 17], to name a few. In classical approaches optical flow is determined by solving a variational optimization problem which requires the judicious choice of a regularization parameter for the problem to be well defined [1, 13, 29]. The choice of regularization parameter turns out critical for the satisfactory inference of optical flow. Despite the many (competing) methods of selecting the regularization parameter, none seems to be most "natural" comparing to the others [12, 24, 28]. Another important feature of classical optical flow approaches is that they produce a single solution (by design) as a "point estimate". In practice, the magnitude of the flow as

2010 *Mathematics Subject Classification.* Primary: 68U10, 62H35; Secondary: 94A08.

Key words and phrases. optical flow, uncertainty quantification, statistical Bayesian inversion.

This work was supported by the National Geospatial-Intelligence Agency under Grant No. HM02101310010 and also by the Office of Naval Research under Grant No. N00014-15-1-2093. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agency.

* Corresponding author: Jie Sun.

well as measurement error and noise can vary significantly from one part of the images to another. Although not captured in classical optical flow, such "uncertainty" would provide valuable information if attainable as part of the solution.

In this work we develop a statistical Bayesian inversion approach for the inference of optical flow. In this framework the optical flow problem is reformulated from the Bayesian perspective as a statistical inversion problem. From this new perspective, different types of information can be naturally fused into a single posterior distribution, which is sampled using an appropriately designed Markov chain Monte Carlo scheme. Unlike the point estimate of optical flow obtained by classical variational approaches, the proposed statistical inversion approach is a methodological way of uncertainty propagation to produce a distribution of candidate optical flow fields from which various statistical properties can be extracted, including ensemble average computation and uncertainty quantification.

The rest of the paper is organized as follows. In Section 2 we develop a discrete approximation of the variational optical flow problem, giving rise to a finite-dimensional linear inverse problem. In Section 3 we reformulate such inverse problem from a Bayesian perspective as a statistical inversion problem and present an efficient algorithm to sample the resulting posterior distribution which is the key to determining the statistical properties of the underlying optical flow fields. In Section 4 we showcase the utility of our Bayesian approach on several benchmark examples and different noise levels. Finally, discussion and conclusion are presented in Section 5.

2. Statistical Inversion Formulation of Optical Flow.

Given a sequence of consecutively captured images, the visual relative motion between them is referred to as optical flow, which can often provide insight about the actual physical motion. The inference of optical flow is an outstanding scientific question, which requires making assumptions about the underlying motion as well as the measurement process.

2.1. Problem Setup.

Consider two single-channeled (typically grayscale) digital images ("pictures") taken from the same scene at two nearby time instances. The image data are represented by two matrices $F = [F_{ij}]_{n_x \times n_y}$ and $G = [G_{ij}]_{n_x \times n_y}$. Thus each image contains $n_x \times n_y$ pixels, defined on a common two-dimensional subspace Ω. The goal is to find matrices $U = [U_{ij}]_{n_x \times n_y}$ and $V = [V_{ij}]_{n_x \times n_y}$ where $\langle U_{ij}, V_{ij} \rangle$ represents the optical "velocity" occurring at the (i,j)-th pixel inferred from the two images. The image data F and G are often regarded as sampled data from smooth functions $F(x,y)$ and $G(x,y)$, with $F_{ij} = F(x_i, y_j)$ and $G_{ij} = G(x_i, y_j)$ where $\{(x_i, y_j)\}_{(i=1,\ldots,n_x; j=1,\ldots,n_y)}$ are grid points from a spatial domain Ω. Thus U and V can be viewed as discrete spatial samples of a smooth velocity field $\vec{W}(x,y) = \langle U(x,y), V(x,y) \rangle$ defined on Ω that captures the visual optical motion occurring between the two observed images.

2.2. Variational Approach of Inferring Optical Flow.

The classical variational approach of optical flow starts by defining an "energy" functional whose minimization yields an estimation of the optical flow field [1]. One of the most widely used functional was proposed by Horn and Schunck in 1981 [13], given by

$$E(U,V) = \iint_{\Omega} (F_x U + F_y V + F_t)^2 dxdy + \alpha \iint_{\Omega} (\|\nabla U\|^2 + \|\nabla V\|^2) dxdy, \qquad (1)$$

where $U(x, y)$ and $V(x, y)$ are smooth functions defined over Ω which represent a candidate flow field. In the Horn-Schunck functional, the first term is often referred to as *data fidelity* as it measures the deviation of the total image intensity from being conservative, that is, the condition

$$dF/dt = F_x U + F_y V + F_t = 0. \tag{2}$$

The second term measures the *solution regularity* by penalizing solutions that have large spatial gradients and is called the regularization term. The relative emphasis of smoothness as compared to "fitting" the total image intensity conservation equation (2) is controlled by the positive scalar α which is called a *regularization parameter*. The main role of the regularization term is to ensure that the minimization of the functional is a well posed problem. Without the regularization term there can typically exists an infinite number of solutions all of which trivially satisfies the conservation equation (2).

Given α, the functions U and V that minimize the Horn-Schunck functional (1) satisfy the Euler-Lagrange equations

$$\begin{cases} F_x(F_x U + F_y V + F_t) = \alpha(U_{xx} + U_{yy}), \\ F_y(F_x U + F_y V + F_t) = \alpha(V_{xx} + V_{yy}), \end{cases} \tag{3}$$

which are typically solved by some iterative scheme over a finite set of spatial grid points [13, 23] to produce an estimation of the optical flow. Alternatively, one could also discretize the functional (1) itself to yield a finite-dimensional inverse problem as discussed in Section 2.3 below with solution strategy reviewed in Section 3.

2.3. **Discrete Approximation of the Variational Optical Flow Functional.** As discussed in the previous section, the classical variational approach of optical flow works by first formulating and minimizing a functional over smooth vector fields, and then evaluating the obtained vector field at the grid points on which the original image data are given. Here we take a different route, by first discretizing the functional (1) to convert the functional minimization (an infinite-dimensional problem) into a finite-dimensional linear inverse problem, and then solving the inverse problem to yield solutions which give the values of a vector field defined over the same grid points as the image data.

The remainder of this section will be focused on the conversion of the functional (1) into a finite-dimensional function defined on a set of uniformly distributed grid points

$$\{(x_i, y_j)\}_{i=1,2,\ldots,n_x; j=1,2,\ldots,n_y}, \tag{4}$$

where $x_{i+1} - x_i = \Delta x$ and $y_{j+1} - y_j = \Delta y$ are the spacing in the x direction and the y direction, respectively. The conversion will be achieved by approximating the integrals in Eq. (1) with appropriately derived summations over the grid points. For notational convenience, we use a bold lowercase variable to denote the vectorization of a matrix. For example, the boldface vector $\boldsymbol{q}$ denotes the column vector obtained by "vertically stacking" the columns of a matrix $Q = [\vec{Q}_1, \vec{Q}_2, \ldots, \vec{Q}_n]$ in order [10], where $\vec{Q}_i$ denotes the i-th column of Q. That is,

$$\boldsymbol{q} = \mathrm{vec}(Q) = \begin{pmatrix} \vec{Q}_1 \\ \vec{Q}_2 \\ \vdots \\ \vec{Q}_n \end{pmatrix}. \tag{5}$$

First let us consider the data fidelity term: $\iint_\Omega (F_x U + F_y V + F_t)^2 dxdy$. The spatial derivatives $F_x(x,y)$ and $F_y(x,y)$ can be approximated by a finite difference scheme. For example, the simple *forward difference* yields the approximations:

$$\begin{cases} F_x(x,y) \approx \frac{1}{\Delta x}\left[F(x+\Delta x, y) - F(x,y)\right], \\ F_y(x,y) \approx \frac{1}{\Delta y}\left[F(x, y+\Delta y) - F(x,y)\right]. \end{cases} \tag{6}$$

We next express these derivatives as operations on the column vector $\boldsymbol{f}$. To do this, we define matrix $S_k = [S_k^{(ij)}]_{k\times k}$ as

$$S_k^{(ij)} = \begin{cases} -\delta_{ij} + \delta_{i+1,j}, & \text{if } i < k; \\ -\delta_{i,j-1} + \delta_{i,j}, & \text{if } i = k. \end{cases} \tag{7}$$

The forward difference applied to $\boldsymbol{f}$ can be represented as

$$\begin{cases} \boldsymbol{f}_x \approx Q_x \boldsymbol{f}, \\ \boldsymbol{f}_y \approx Q_y \boldsymbol{f}, \end{cases} \tag{8}$$

where

$$\begin{cases} Q_x \equiv \frac{1}{\Delta x}[I_n \otimes S_m], \\ Q_y \equiv \frac{1}{\Delta y}[S_n \otimes I_m]. \end{cases} \tag{9}$$

The temporal derivative can be estimated from the data by the difference between the two images, to yield

$$\boldsymbol{f}_t \approx \boldsymbol{g} - \boldsymbol{f}. \tag{10}$$

With these definitions, we obtain a discretized version of the conservation equation (2) expressed as a linear system:

$$A\boldsymbol{x} = \boldsymbol{b}, \tag{11}$$

where

$$\begin{cases} A = [\mathrm{diag}(\boldsymbol{f}_x), \mathrm{diag}(\boldsymbol{f}_y)], \\ \boldsymbol{x} = [\boldsymbol{u}^\top, \boldsymbol{v}^\top]^\top, \\ \boldsymbol{b} = -\boldsymbol{f}_t. \end{cases} \tag{12}$$

Here $\mathrm{diag}(\boldsymbol{f})$ represents a diagonal matrix whose diagonal elements are given by the entries of the vector $\boldsymbol{f}$. From this connection we approximate the first integral in the functional (1) as $\|A\boldsymbol{x} - \boldsymbol{b}\|^2$ where $\|\cdot\|$ denotes the standard Euclidean norm.

Next we develop a finite-dimensional approximation of the regularization term in the functional (1). This requries discretization of ∇U and ∇V. Using a similar forward difference to approximate the partial derivatives, we obtain

$$\begin{cases} \nabla U \approx \frac{1}{\Delta x}\left[U(x+\Delta x, y) - U(x,y)\right] + \frac{1}{\Delta y}\left[U(x, y+\Delta y) - U(x,y)\right], \\ \nabla V \approx \frac{1}{\Delta x}\left[V(x+\Delta x, y) - V(x,y)\right] + \frac{1}{\Delta y}\left[V(x, y+\Delta y) - V(x,y)\right]. \end{cases} \tag{13}$$

For the vectorized variables $\boldsymbol{u}$ and $\boldsymbol{v}$, we have

$$\begin{cases} \nabla \boldsymbol{u} = \boldsymbol{u}_x + \boldsymbol{u}_y \approx (Q_x + Q_y)\boldsymbol{u}, \\ \nabla \boldsymbol{v} = \boldsymbol{v}_x + \boldsymbol{v}_y \approx (Q_x + Q_y)\boldsymbol{v}, \end{cases} \tag{14}$$

where Q_x and Q_y are defined in Eq. (9). Consequently, we obtain the approximation of the second integral in the Horn-Schunck functional (1) as

$$\begin{cases} \iint_\Omega \|\nabla U\|^2 dxdy \approx \boldsymbol{u}_x^\top \boldsymbol{u}_x + \boldsymbol{u}_y^\top \boldsymbol{u}_y \approx \boldsymbol{u}^\top [Q_x^\top Q_x + Q_y^\top Q_y]\boldsymbol{u}, \\ \iint_\Omega \|\nabla V\|^2 dxdy \approx \boldsymbol{v}_x^\top \boldsymbol{u}_y + \boldsymbol{v}_y^\top \boldsymbol{v}_y \approx \boldsymbol{v}^\top [Q_x^\top Q_x + Q_y^\top Q_y]\boldsymbol{v}, \end{cases} \tag{15}$$

which then gives

$$\iint_{\Omega} (\|\nabla u\|^2 + \|\nabla v\|^2) dxdy \approx \boldsymbol{x}^\top Q \boldsymbol{x}, \tag{16}$$

where the matrix

$$Q = I_2 \otimes [Q_x^\top Q_x + Q_y^\top Q_y]. \tag{17}$$

Therefore, the variational optical flow formulation (1) can be reformulated at a finite spatial resolution as the following regularized optimization problem:

$$\min_{\boldsymbol{x}} \left(\|A\boldsymbol{x} - \boldsymbol{b}\|^2 + \alpha \boldsymbol{x}^\top Q \boldsymbol{x} \right), \tag{18}$$

which is a standard linear least squares problem with Tikhonov regularization, with more details to be presented in the next section.

3. Inverse Problem in Finite Dimensions.

Although there has been a great deal of progress on the mathematical characterization of inverse problems in the terms of functional analysis, a practical problem often concerns finding a solution in a finite-dimensional space. At a fundamental level, the most common inverse problem stems from a linear model [12, 14, 28]

$$\boldsymbol{b} = A\boldsymbol{x} + \boldsymbol{\eta}, \tag{19}$$

where $\boldsymbol{b} \in \mathbb{R}^m$ is a column vector of *observed data*, $A = [a_{ij}]_{m\times n} \in \mathbb{R}^{m\times n}$ is a (known) matrix representing the underlying model, the column vector $\boldsymbol{\eta} \in \mathbb{R}^n$ denotes (additive) noise, and $\boldsymbol{x} \in \mathbb{R}^n$ is the vector of *unknowns* to be inferred.

Given A and $\boldsymbol{b}$, the problem of inferring $\boldsymbol{x}$ in Eq. (19) is called an *inverse problem* because rather than direct "forward" computation from the model, it requires a set of indirect, "backward", or "inverse" operations to determine the unknowns [28]. Depending on the rank and conditioning of the matrix A, the problem may be ill-posed or ill-conditioned. In classical approaches, these issues are dealt with by adjusting the original problem to a (slightly) modified optimization problem as discussed in Section 3.1 whose solution is meant to represent the original, as discussed in Section 3.2.

We note that in the classical setting a solution to the inverse problem is a vector $\boldsymbol{x}$ as a result of solving an optimization problem. Such a solution is referred to as an *point estimate* because it gives one solution vector without providing any information about how reliable (or uncertain) the solution is [12, 28]. On the other hand, the statistical inversion approach to inverse problems provides an *ensemble* of solutions together with a distribution from which not only point estimates can be made but also their uncertainty quantification [9, 14].

3.1. Least squares solution.

The classical least squares solution to the inverse problem is given by [10]

$$\boldsymbol{x}_{\ell_2} = A^\dagger \boldsymbol{b}, \tag{20}$$

where $A^\dagger$ denotes the pseudo-inverse of A which can be obtained from the singular value decomposition of A [10]. Depending on the rank of A, the least squares solution $\boldsymbol{x}_{\ell_2}$ is associated with one of the minimization problems:

$$\begin{cases} \min_{A\boldsymbol{x}=\boldsymbol{b}} \|\boldsymbol{x}\|_2, & \text{if } \operatorname{rank}(A) < n \text{ (which is necessarily the case if } m < n); \\ \min_{\boldsymbol{x}} \|A\boldsymbol{x} - \boldsymbol{b}\|_2, & \text{if } \operatorname{rank}(A) = n. \end{cases} \tag{21}$$

Here $\|\cdot\|_2$ denotes the ℓ_2 (Euclidean) norm. Let the true solution to Eq. (19) be $\boldsymbol{x}^*$, that is, $\boldsymbol{b} = A\boldsymbol{x}^* + \boldsymbol{\eta}$. It follows that

$$\boldsymbol{x}_{\ell_2} - \boldsymbol{x}^* = A^\dagger \boldsymbol{\eta}. \tag{22}$$

In practice, even when the matrix A has full column rank ($\text{rank}(A) = n$), the discrepancy between the true and least squares solutions is typically dominated by noise when some singular values of A are close to zero, rendering A an ill-conditioned matrix and the solution $\boldsymbol{x}^*$ unstable [28].

3.2. **Tikhonov Regularization.** A powerful approach to resolve the instabilities due to noise and the near-singularity of A is to *regularize* the problem. In the classical Tikhonov regularization one adds a quadratic regularization term $\alpha \boldsymbol{x}^T L \boldsymbol{x}$ to penalize non-smoothness, giving rise to a regularized problem [25, 26, 27, 28]:

$$\min_{\boldsymbol{x}} \left(\|A\boldsymbol{x} - \boldsymbol{b}\|_2^2 + \alpha \boldsymbol{x}^T L \boldsymbol{x} \right). \tag{23}$$

In the regularized problem the positive scalar parameter α controls the weight of regularization and L is typically a symmetric positive definite matrix, both of which need to be chosen appropriately for the problem to be uniquely defined [25, 26, 27, 28]. We refer to the two terms $\|A\boldsymbol{x} - \boldsymbol{b}\|_2^2$ and $\alpha \boldsymbol{x}^\top L \boldsymbol{x}$ in (23) as *data fidelity* and *solution regularity*, respectively. In a simplistic description, they can be described as "selecting" a solution $\boldsymbol{x}_\alpha$ that balances the desire to "solve" $A\boldsymbol{x} = \boldsymbol{b}$ and to be "regular" as measured by $\boldsymbol{x}^\top L \boldsymbol{x}$. The regularization parameter α therefore dictates the extent to which the compromise is made between the two.

For a fixed α, we denote the corresponding regularized solution by

$$\boldsymbol{x}_\alpha = \text{argmin}_{\boldsymbol{x}} \{ \|A\boldsymbol{x} - \boldsymbol{b}\|^2 + \alpha \boldsymbol{x}^\top L \boldsymbol{x} \}. \tag{24}$$

By standard vector calculus, it can be shown that $\boldsymbol{x}_\alpha$ is in fact a solution to the modified linear system

$$(A^\top A + \alpha L)\boldsymbol{x}_\alpha = A^\top \boldsymbol{b}, \tag{25}$$

which is typically well-posed for appropriate choices of L and α. When the matrices are large and sparse, Eq. (25) can be solved by iterative methods rather than a direct matrix inversion which tends to be numerically costly and unstable [10].

The key remains how to appropriately choose the regularization parameter α. Despite the existence and ongoing development of many competing methods for selecting α most of which focus on asymptotical optimality as the number of data points approach infinity, none of them stands out as a "natural" choice unless specific priori information about the noise in the data are available (see Chapter 7 of Ref. [28]). In the following section we show that this problem of selecting an appropriate regulation parameter can be effectively bypassed under a statistical inversion framework from a Bayesian perspective.

4. **Statistical Inversion Approach.** The statistical inversion approach to an inverse problem starts with using the Bayes rule to express the *posterior* distribution $p(\boldsymbol{x}|\boldsymbol{b})$, which is the conditional distribution of the "solution vector" $\boldsymbol{x}$ given the observed data $\boldsymbol{b}$, as [9, 14]

$$p(\boldsymbol{x}|\boldsymbol{b}) = \frac{1}{p(\boldsymbol{b})} p(\boldsymbol{b}|\boldsymbol{x}) \cdot p(\boldsymbol{x}). \tag{26}$$

Here the *likelihood* function $p(\boldsymbol{b}|\boldsymbol{x})$ is the probability density function (pdf) of the random variable $\boldsymbol{b}$ given $\boldsymbol{x}$ which is determined by the underlying model; $p(\boldsymbol{x})$ is

the *priori* distribution of $\boldsymbol{x}$; and $p(\boldsymbol{b}) > 0$ acts as a normalization constant which does not affect the solution procedure or the final solution itself.

The key of the statistical inversion formulation is that each candidate solution $\boldsymbol{x}$ is associated with the probability $p(\boldsymbol{x}|\boldsymbol{b})$ that is determined (up to a normalization constant $1/p(\boldsymbol{b})$) once the likelihood function and the prior distribution are given. For a given inverse problem, the likelihood function is determined by the underlying model such as Eq. (19) including the noise distribution. On the other hand, the prior distribution $p(\boldsymbol{x})$ is typically constructed according to some prior knowledge of the solution. To obtain desired solutions as well as their statistical properties such as uncertainty quantification, one typically needs to *sample* from the posterior distribution. Efficient sampling methods will be reviewed toward the end of this section.

The unique feature of enabling information fusion and uncertainty quantification has made the statistical inversion approach to inverse problems an attractive venue for the development of new theory and applications. In image processing applications, it has been utilized for many problems such as image denoising and deblurring [2, 15], sparse signal reconstruction [19], and more recently for optical flow computation [7]. In particular, we note that our approach, although different in many of the technical aspects, shares a similar statistical inversion perspective as Ref. [7].

4.1. **Tikhonov regularization and MAP solution.** Under the statistical inversion framework, a particularly popular solution is one that maximizes the posteriori probability given the measured data. Such a solution $\boldsymbol{x}_{\text{MAP}}$, called a *maximum a posteriori* (MAP) estimator, is defined as

$$\begin{aligned} \boldsymbol{x}_{\text{MAP}} &= \text{argmax}_{\boldsymbol{x}} p(\boldsymbol{x}|\boldsymbol{p}) = \text{argmin}_{\boldsymbol{x}}\{-\ln p(\boldsymbol{x}|\boldsymbol{b})\} && (27)\\ &= \text{argmin}_{\boldsymbol{x}}\{-\ln p(\boldsymbol{b}|\boldsymbol{x}) - \ln p(\boldsymbol{x})\}. && (28) \end{aligned}$$

As noted in Refs. [2, 14], the MAP estimation given by Eq. (28) is in fact equivalent to the Tikhonov regularization specified in Eq. (24) upon appropriate choice of the model and prior pdf. In particular, consider the model given by Eq. (19) with independent and identically distributed (iid) Gaussian noise of variance λ^{-1}. It follows that

$$p(\boldsymbol{b}|\boldsymbol{x}) = p(\boldsymbol{\eta}) \propto \exp\left(-\frac{\lambda}{2}\|A\boldsymbol{x} - \boldsymbol{b}\|^2\right), \tag{29}$$

where the symbol "$\propto$" means "proportional to". If the prior is also taken to be a Gaussian distribution as

$$p(\boldsymbol{x}) \propto \exp\left(-\frac{\delta}{2}\boldsymbol{x}^\top L\boldsymbol{x}\right), \tag{30}$$

then the term $-\ln p(\boldsymbol{x}|\boldsymbol{b})$ in the MAP estimator becomes

$$-\ln p(\boldsymbol{x}|\boldsymbol{b}) \propto \|A\boldsymbol{x} - \boldsymbol{b}\|^2 + (\delta/\lambda)\boldsymbol{x}^\top L\boldsymbol{x}. \tag{31}$$

The choice of $\delta/\lambda = \alpha$ then yields the same MAP estimator $\boldsymbol{x}_{\text{MAP}}$ as the regularization solution $\boldsymbol{x}_\alpha$ given by Eq. (24). Thus, with these assumptions of the form of the noise, the distribution of the prior, there is a logical bridge between two different philosophies for inverse problems, since the MAP estimator from the posterior distribution under the statistical inversion framework is equivalent to the result of a specific Tikhonov regularization solution.

Furthermore, and crucially, as shown in the next section, more important information exists in the statistical inversion framework. Specifically, by sampling from the posterior distribution the statistical inversion approach allows for not only a point estimate but also other statistical properties associated with the solution, in particular uncertainty quantification.

4.2. **Computational Aspects.** In the statistical inversion formalism, once the form of the posterior distribution is derived, the remaining part of the work is devoted to efficient sampling from the posterior distribution. Typically a Markov chain Monte Carlo (MCMC) sampling approach is adopted. The main idea is to generate a sequence of samples according to a prescribed Markov chain whose unique stationary distribution is the desired posterior distribution.

In this paper we will follow the work by [2] to consider a specific class of the noise and prior distributions:

$$\begin{cases} \text{likelihood function: } p(\boldsymbol{b}|\boldsymbol{x},\lambda) \propto \lambda^{m/2} \exp\left(-\frac{\lambda}{2}\|A\boldsymbol{x}-\boldsymbol{b}\|^2\right), \\ \text{prior distribution: } p(\boldsymbol{x}|\delta) \propto \delta^{n/2} \exp\left(-\frac{\delta}{2}\boldsymbol{x}^\top L\boldsymbol{x}\right), \end{cases} \tag{32}$$

Here the noise is assumed to be additive, Gaussian, and independent of the measured data, with variance λ^{-1}, giving rise to the form of the likelihood function. On the other hand, the prior distribution is considered to be Gaussian with covariance matrix $(\delta L)^{-1}$ (matrix δL is referred to as the *precision* matrix). For optical flow applications, we will choose $L = Q$ where Q is given by Eq. (17) which corresponds to a spatial regularization measure. As it turns out, this choice of L is closely related to the selection of prior according to a spatial Gaussian Markov random field which is common in tackling spatial inverse problems [3, 11].

To completely specify the posterior distribution, we also need to choose prior distributions for the parameters λ and δ. These are often called *hyperpriors*. Following Ref. [2], we choose the priors for $p(\lambda)$ and $p(\delta)$ to be Gamma distributions, as

$$\begin{cases} p(\lambda) \propto \lambda^{\alpha_\lambda - 1} \exp(-\beta_\lambda \lambda), \\ p(\delta) \propto \delta^{\alpha_\delta - 1} \exp(-\beta_\delta \delta). \end{cases} \tag{33}$$

Such choice ensures that $p(\lambda)$ and $p(\delta)$ are conjugate hyper-priors. Typically, without much prior knowledge of the values of λ and δ one would choose the values of α_λ, β_λ, α_δ, and β_δ to ensure the distributions $p(\lambda)$ and $p(\delta)$ to be "wide", allowing the Markov chain to have the opportunity to explore a large part of the parameter space. Following the remedy suggested in Ref. [2, 3, 11], we set $\alpha_\lambda = \alpha_\delta = 1$ and $\beta_\lambda = \beta_\delta = 10^{-4}$ unless otherwise noted. We tested other choice of parameters as well and they mainly affect the length of the transient in the MCMC sampling process and do not seem to have a strong influence on the asymptotic outcome.

Consequently, the full conditional distributions that relate to the posterior distribution are given by

$$\begin{cases} p(\boldsymbol{x}|\lambda,\delta,\boldsymbol{b}) \propto \exp\left(-\frac{\lambda}{2}\|A\boldsymbol{x}-\boldsymbol{b}\|^2 - \frac{\delta}{2}\boldsymbol{x}^\top L\boldsymbol{x}\right), \\ p(\lambda|\boldsymbol{x},\delta,\boldsymbol{b}) \propto \lambda^{m/2+\alpha_\lambda-1} \exp\left(\lambda\left[-\frac{1}{2}\|A\boldsymbol{x}-\boldsymbol{b}\|^2 + \beta_\lambda\right]\right), \\ p(\delta|\boldsymbol{x},\lambda,\boldsymbol{b}) \propto \delta^{n/2+\alpha_\delta-1} \exp\left(\delta\left[-\frac{1}{2}\boldsymbol{x}^\top L\boldsymbol{x} - \beta_\delta\right]\right). \end{cases} \tag{34}$$

In other words,

$$\begin{cases} \boldsymbol{x}|\lambda,\delta,\boldsymbol{b} \sim \mathcal{N}\left((\lambda A^\top A+\delta L)^{-1}\lambda A^\top \boldsymbol{b}, (\lambda A^\top A+\delta L)^{-1}\right), \\ \lambda|\boldsymbol{x},\delta,\boldsymbol{b} \sim \Gamma\left(m/2+\alpha_\lambda, \frac{1}{2}\|A\boldsymbol{x}-\boldsymbol{b}\|^2+\beta_\lambda\right), \\ \delta|\boldsymbol{x},\lambda,\boldsymbol{b} \sim \Gamma\left(n/2+\alpha_\delta, \frac{1}{2}\boldsymbol{x}^\top L\boldsymbol{x}+\beta_\delta\right). \end{cases} \tag{35}$$

Here we emphasize that for applications with non-square matrix $A_{m\times n}$ (such as our optical flow application), it is important to have m and n in the right places to ensure an appropriately defined MCMC.

4.3. **A Gibbs sampler of the posterior distribution.** We adopt the block Gibbs sampler developed in Refs. [2, 14] as a specific MCMC procedure to sample the posterior distribution. In theory the sample distribution asymptotically converges to the true posterior distribution. The approach contains the following steps.

0. Initialize δ_0 and λ_0, and set $k=0$.
1. Sample $\boldsymbol{x}^k \sim \mathcal{N}\left((\lambda A^\top A+\delta L)^{-1}\lambda A^\top \boldsymbol{b}, (\lambda A^\top A+\delta L)^{-1}\right)$.
2. Sample $\lambda_{k+1} \sim \Gamma\left(m/2+\alpha_\lambda, \frac{1}{2}\|A\boldsymbol{x}^k-\boldsymbol{b}\|^2+\beta_\lambda\right)$.
3. Sample $\delta_{k+1} \sim \Gamma\left(n/2+\alpha_\delta, \frac{1}{2}(\boldsymbol{x}^k)^\top L(\boldsymbol{x}^k)+\beta_\delta\right)$.
4. Set $k \leftarrow k+1$ and return to Step 1.

Here the computational burden is mainly due to Step 1, which requires drawing samples from a multivariate Gaussian variable, which is equivalent to solving the following linear system at each iteration for $\boldsymbol{x}^k$:

$$(\lambda_k A^\top A+\delta_k L)\boldsymbol{x}^k = \lambda_k A^\top \boldsymbol{b}+\boldsymbol{w}, \text{ where } \boldsymbol{w} \sim \mathcal{N}(\boldsymbol{0}, \lambda_k A^\top A+\delta_k L). \tag{36}$$

For large matrices, instead of a direct solve using Gauss elimination, an iterative method is usually preferred. Among the various notable iterative methods such as Jacobi, Gauss-Seidel (G-S), and conjugate gradient (CG) [18], we adopted the CG for all the numerical experiments as reported in this paper, with a starting vector of all zeros, maximum of 500 iterations, and error tolerance of 10^{-6}

5. **Examples of Statistical Inversion Based Optical Flow.** To test our proposed statistical inversion approach to optical flow, we consider several synthetic image pairs for which the ground truth optical flow (U,V) is known. In each example, we generate two images according to the following equation

$$\boldsymbol{g} = \boldsymbol{f} - \boldsymbol{f}_x\boldsymbol{u} - \boldsymbol{f}_y\boldsymbol{v} + \boldsymbol{\eta}, \tag{37}$$

where $\boldsymbol{f}$ is the first image, $\boldsymbol{g}$ is the second image, and $\boldsymbol{\eta}$ denotes (multivariate) noise whose individual components are independently drawn from a Gaussian distribution with zero mean and fixed standard deviation σ.

We consider 5 qualitatively different optical flow fields, all defined on the same normalized spatial domain $[-1,1]\times[-1,1]$, with 30 uniform grid points in each direction, resulting in image matrices of size 30-by-30. The first image in all the examples are identical, generated by the equation

$$F(x,y) = \frac{1}{2}\left[\cos(\pi x)\cos(\pi y)+1\right], \tag{38}$$

as shown in panel (a) of Figures 1 to 5. The second image G is generated according to Eq. (37) using various optical flow fields and noise levels as described below, with the spatial derivatives numerically implemented using the forward difference scheme.

Example 1: $\langle U(x,y), V(x,y)\rangle = \langle x, y\rangle$.

Example 2: $\langle U(x,y), V(x,y)\rangle = \langle -y, x\rangle$.

Example 3: $\langle U(x,y), V(x,y)\rangle = \langle y, \sin(x)\rangle$.

Example 4: $\langle U(x,y), V(x,y)\rangle = \langle -\pi \sin(0.5\pi x)\cos(0.5\pi y), \pi\cos(0.5\pi x)\sin(0.5\pi y)\rangle$.

Example 5: $\langle U(x,y), V(x,y)\rangle = \langle -\pi \sin(\pi x)\cos(\pi y), \pi\cos(\pi x)\sin(\pi y)\rangle$.

For each optical flow field, we generate the second image G either without noise or with noise (standard deviation $\sigma = 0.02$). Then, we adopt MCMC-Gibbs sampling procedure and corresponding choice of prior pdf and hyerpriors presented in Section 4 to produce a posterior distribution $p(U, V)$ of the optical flow field. Having a full distribution (the posterior distribution) rather than a single solution is a unique feature of the Bayesian approach to optical flow, which enables one to enumerate and quantify uncertainty as desired.

Figures 1∼5 show the results of the statistical inversion based optical flow for the 5 test example, respectively. In each figure, the top row (a1-a3) shows the image data of the first image F (a1), and the second image G generated from Eq. (37) with no noise (a2) and with noise under standard deviation $\sigma = 0.02$ (a3), respectively. The middle rows (b1-b3) show the true optical flow field (b1) compared with the inferred "mean" optical flow fields together with uncertainty quantification from the MCMC samples (b2-b3). In particular, at each point $z = (x, y)$, we construct a 2d normal pdf $\mathcal{N}(\mu, \Sigma)$ by using the sample mean μ and sample covariance Σ estimated from the MCMC samples after discarding the initial transients. This allows us to obtain a "mean" optical flow at point z defined as $\langle u(z), v(z)\rangle = \langle \mu_1, \mu_2\rangle$. Uncertainty is quantified by computing a confidence region that contains q probability mass of the fitted multivariate normal distribution given by [20]

$$(z-\mu)^\top \Sigma^{-1}(z-\mu) \le \chi_2^2(q). \tag{39}$$

Here $\chi_2^2(q)$ denotes the q-th quantile of the Chi-squared distribution with two degrees of freedom, that is, $\chi_2^2(q) = K^{-1}(q)$ where K is the cdf of χ_2^2. These confidence regions (shaded ellipses) are shown for the zoomed-in plots for the inferred optical flow fields. Finally, the last row (c1-c3) in each figure shows how the MCMC procedure produces a distribution of the effective regularization parameter σ/λ. Panel (c1) shows the change of σ/λ over time in the MCMC sampling procedure, indicating convergence to a stationary distribution typically after a quick initial transient. The remaining panels (c2) and (c3) show the distribution of σ/λ after discarding the initial transient, for both the case of no noise (c2) and the case with noise (c3).

We point out a few observations from the numerical experiments. First, the estimated optical flow compare reasonably well with the true flow in all examples of the qualitatively different optical flow fields, supporting the utility of the proposed statistical inversion approach [see panels (b1-b3) in all figures]. Secondly, we again point out that the MCMC procedure used in our statistical inversion approach to optical flow does not require an active prior choice of the regularization parameter. The MCMC samples seems to quickly converge to a stationary distribution for the effective parameter [panel (c1) in all figures], from which the distributions of parameters and solutions can be determined. Finally, comparing to the noise free images, the estimation of optical flow becomes less accurate when noise is added. It is worth mentioning that the statistical inversion approach in fact allows us to "predict" this difference without knowing the ground-truth optical flow, by quantifying and comparing the uncertainty of solutions [panels (b2) and (b3) in all figures].

6. **Discussions and Conclusion.** In this paper we take a statistical inversion approach to the optical flow inference problem. The key step is to formulate a linear inverse problem from which the posterior distribution can be expressed by utilizing knowledge about the form of model, noise, and other prior information. From a Bayesian perspective, this information is combined to produce a posterior distribution describing the propagation of prior information in context of the problem. We have shown that traditional variational approaches such as the seminal work developed by Horn and Schunck [13] can in fact be regarded as a special case within the statistical inversion framework by making specific assumptions about the model, noise, and prior distribution. Thus we recap that there are major advantages over the classical variational calculus approach to inverse problems where by necessity the ill-posedness is dealt with by adding an ad hoc regularity term that hopefully agrees with expected physical interpretation. From a Bayesian perspective, the ill-posedness is dealt with naturally under the statistical inversion framework by restating as a well-posed extended problem in a larger space of probability distributions [9, 14]. This therefore naturally removes a key difficulty of choosing an appropriate regularity parameter encountered in classical methods. Instead, in contrast to classical optical flow methods which only yield single solutions as "point estimates", the statistical inversion approach produces a distribution of solutions which can be sampled in terms of most appropriate estimators and also for uncertainty quantification. Specifically in the context of an optic flow problem, we expect a distribution of regularity parameters, and correspondingly a distribution of optical flow vectors at each point.

In this paper we focused on a statistical inversion formulation that is based on the classical Horn-Schunck functional, other functionals can be similarly coped with so long as a linear inverse problem can be formulated. The Horn-Schunck framework assumed rigid body motion and conservation of brightness, giving rise to a divergence-free functional. Other data fidelity terms [5, 16, 6] can be formulated to correspond to the physics of the underlying application, for example for fluid and oceanographic problems where a stream function, or even a quasi-static approximation to assume such physics as corriolis can be used. Likewise, regularity in time and multiple time shots may be appropriate [5], as these correspondingly more complex formulations nonetheless come back to a linear inverse problem tenable in the framework of this paper. Specifically within the statistical inversion framework developed in the current paper, all of these could be recast so that the data fidelity term may be written into the basic form here to allow the statistical inverse problems framework, which we plan in future work. Likewise, other numerical differentiation and integration schemes can be used as well in place of the simple forward difference used here.

Acknowledgments. We thank Ranil Basnayake for many insightful discussions and constructive comments on the manuscript.

REFERENCES

[1] G. Aubert, R. Deriche and P. Kornprobst, Computing optical flow via variational techniques, *SIAM J. Appl. Math.*, **60** (1999), 156–182.

[2] J. M. Bardsley, MCMC-based image reconstruction with uncertainty quantification, *SIAM J. Sci. Comput.*, **34** (2012), A1316–A1332.

[3] J. M. Bardsley, Gaussian Markov random field priors for inverse problems, *Inverse Problems and Imaging*, **7** (2013), 397–416.

[4] J. L. Barron, D. J. Fleet and S. S. Beauchemin, Performance of optical flow techniques, *International Journal of Computer Vision*, **12** (1994), 43–77.
[5] R. Basnayake and E. M. Bollt, A Multi-Time Step Method to Compute Optical Flow with Scientific Priors for Observations of a Fluidic System, in *Ergodic Theory, Open Dynamics, and Coherent Structures* (eds. G. Froyland and W. Bahsoun), Springer, 2014.
[6] R. Basnayake, A. Luttman and E. M. Bollt, A lagged diffusivity method for computing total variation regularized fluid flow, *Contemp. Math*, **586** (2013), 57–64.
[7] G. Chantas, T. Gkamas and C. Nikou, Variational-Bayes optical flow, *J. Math. Imaging Vision*, **50** (2014), 199–213.
[8] I. Cohen and I. Herlin, Optical flow and phase portrait methods for environmental satellite image sequences, in *Proc. Europ. Conf. Computer Vision*, Cambridge, U.K., 1996, pp. 141–150.
[9] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari and D. B. Rubin, *Bayesian Data Analysis*, 3^{rd} edition, Chapman & Hall/CRC, 2014.
[10] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 4^{th} edition, Johns Hopkins University Press, 2012.
[11] D. Higdon, A primer on space-time modelling from a Bayesian perspective, in *Los Alamos National Laboratory, Statistical Sciences Group, Technical Report*, LA-UR-05-3097.
[12] P. C. Hansen, *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*, SIAM, Philadelphia, 1998.
[13] B. K. P. Horn and B. G. Schunck, Determining optical flow, *Artificial Intelligence*, **17** (1981), 185–203.
[14] J. Kaipio and E. Somersalo, *Statistical and Computational Inverse Problems*, Springer, 2005.
[15] M. Lebrun, A. Buades and J. M. Morel, A nonlocal Bayesian image denoising algorithm, *SIAM J. Imaging Sci.*, **6** (2013), 1665–1688.
[16] A. Luttman, E. Bollt, R. Basnayake and S. Kramer, A stream function approach to optical flow with applications to fluid transport dynamics, *Proc. Appl. Math. Mech.*, **11** (2011), 855–856.
[17] A. Luttman, E. M. Bollt, R. Basnayake, S. Kramer and N. B. Tufillaro, A framework for estimating potential fluid flow from digital imagery, *Chaos*, **23** (2013), 033134.
[18] Y. Saad, *Iterative methods for sparse linear systems*, 2^{nd} edition, SIAM, 2003.
[19] M. W. Seeger and H. Nickisch, Large scale Bayesian inference and experimental design for sparse linear models, *SIAM J. Imaging Sci.*, **1** (2011), 166–199.
[20] M. Siotani, Tolerance regions for a multivariate normal population, *Annals of the Institute of Statistical Mathematics*, **16** (1964), 135–153.
[21] K. Souhila and A. Karim, Optical flow based robot obstacle avoidance, *International Journal of Advanced Robotic Systems*, **4** (2007), 13–16.
[22] R. Szeliski, *Computer Vision*, Springer 2011.
[23] L. Le Tarnec, F. Destrempes, G. Cloutier and D. Garcia, A proof of convergence of the Horn-Schunck optical flow algorithm in arbitrary dimension, *SIAM J. Imaging Sci.*, **7** (2014), 277–293.
[24] A. M. Thompson, J. C. Brown, J. W. Kay and D. M. Titterington, A study of methods of choosing the smoothing parameter in image restoration by regularization, *IEEE Trans. Pattern Anal. Mach. Intell.*, **13** (1991), 326–339.
[25] A. N. Tikhonov, Regularization of incorrectly posed problems, *Soviet Mathematics Doklady*, **4** (1963), 1624–1627.
[26] A. N. Tikhonov and V. Arsenin, *Solutions of Ill-Posed Problems*, Wiley, New York, 1977.
[27] A. N. Tikhonov, A. V. Goncharsky, V. V. Stepanov and A. G. Yagola, *Numerical Methods for the Solution of Ill-Posed Problems*, Kluwer Academic Publishers, 1990.
[28] C. R. Vogel, *Computational Methods for Inverse Problems*, SIAM, Philadelphia, 2002.
[29] J. Weickert and C. Schnörr, Variational optic flow computation with a spatio-temporal smoothness constraint, *Journal of Mathematical Imaging and Vision*, **14** (2001), 245–255.

Received xxxx 20xx; revised xxxx 20xx.

E-mail address: sunj@clarkson.edu
E-mail address: bolltem@clarkson.edu

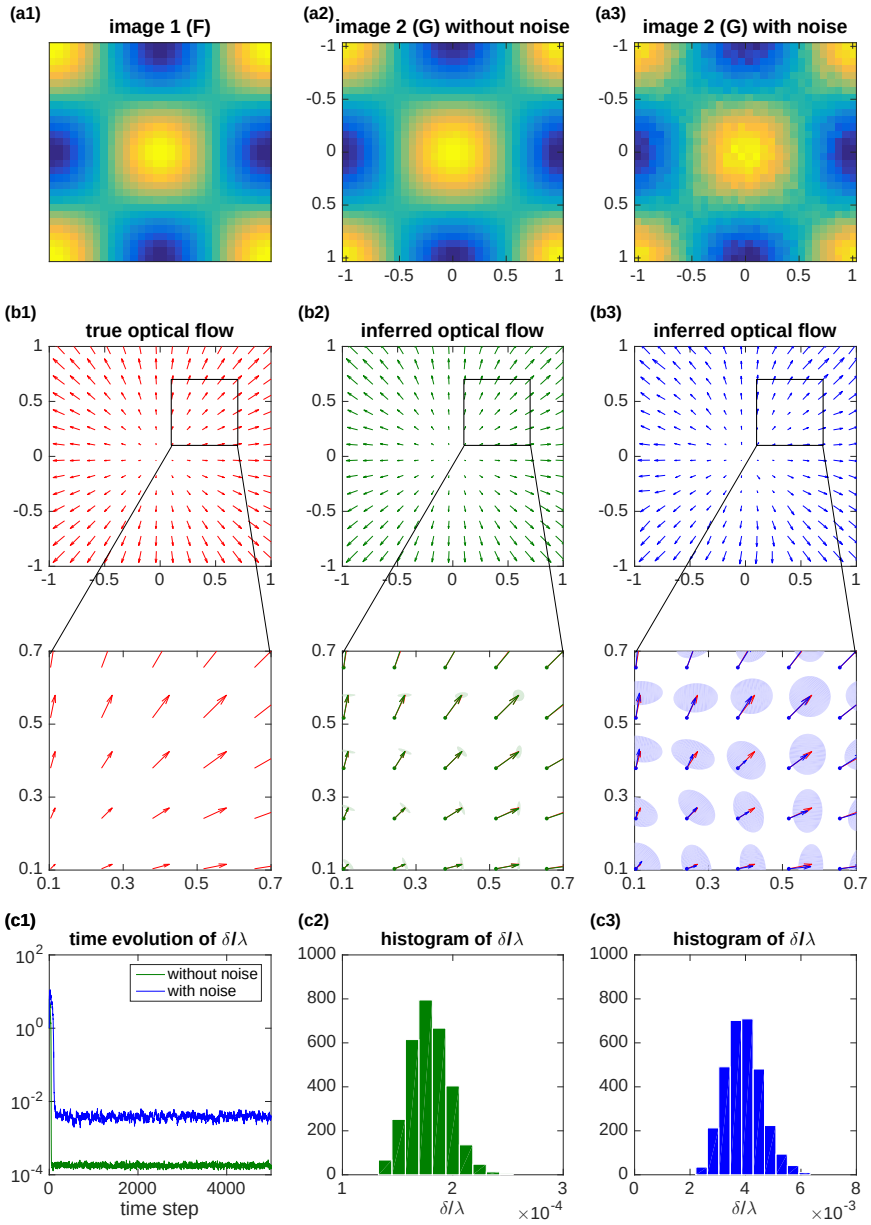

FIGURE 1. Statistical inversion based optical flow example 1: $\langle U(x,y), V(x,y)\rangle = \langle x, y\rangle$. Top row (a1-a3): image data generated according to Eq. (37) for F (a1) and using Eq. (38) for G either without noise (a2) or with noise (a3). Middle rows (b1-b3): the true optical flow field (b1) and the inferred optical flow field from the MCMC samples. In all panels corresponding to the inferred optical flow, the mean vector at each point is plotted as an arrow; for the zoomed-in views, we show the computed confidence region defined by Eq. (39) as shaded ellipses around the tip of the arrows. Bottom row (c1-c3): time evolution (c1) as well as the distribution (c2-c3) of the effective regularization parameter σ/λ, where the distributions are obtained after discarding the initial transient in the MCMC sampling process.

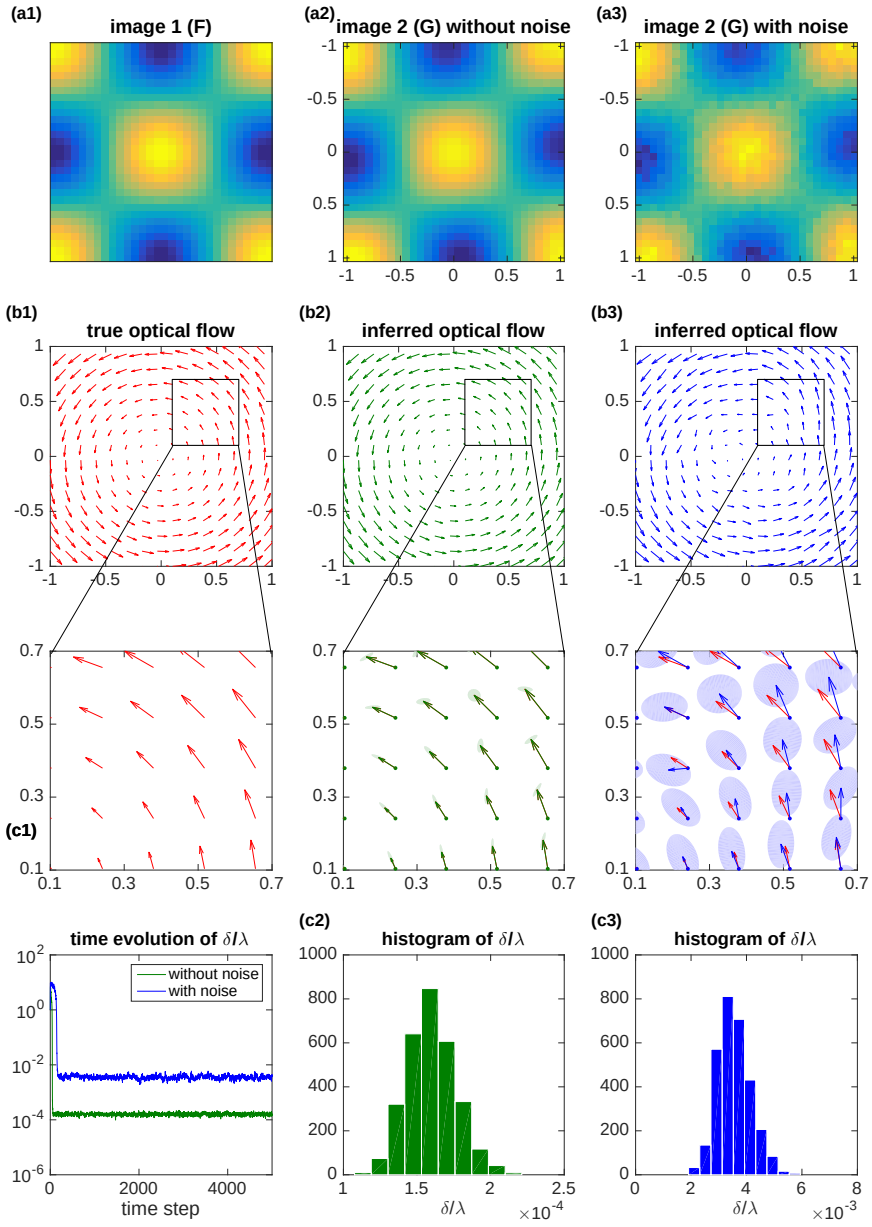

FIGURE 2. Statistical inversion based optical flow example 1: $\langle U(x,y), V(x,y)\rangle = \langle x, y\rangle$. Top row (a1-a3): image data generated according to Eq. (37) for F (a1) and using Eq. (38) for G either without noise (a2) or with noise (a3). Middle rows (b1-b3): the true optical flow field (b1) and the inferred optical flow field from the MCMC samples. In all panels corresponding to the inferred optical flow, the mean vector at each point is plotted as an arrow; for the zoomed-in views, we show the computed confidence region defined by Eq. (39) as shaded ellipses around the tip of the arrows. Bottom row (c1-c3): time evolution (c1) as well as the distribution (c2-c3) of the effective regularization parameter σ/λ, where the distributions are obtained after discarding the initial transient in the MCMC sampling process.

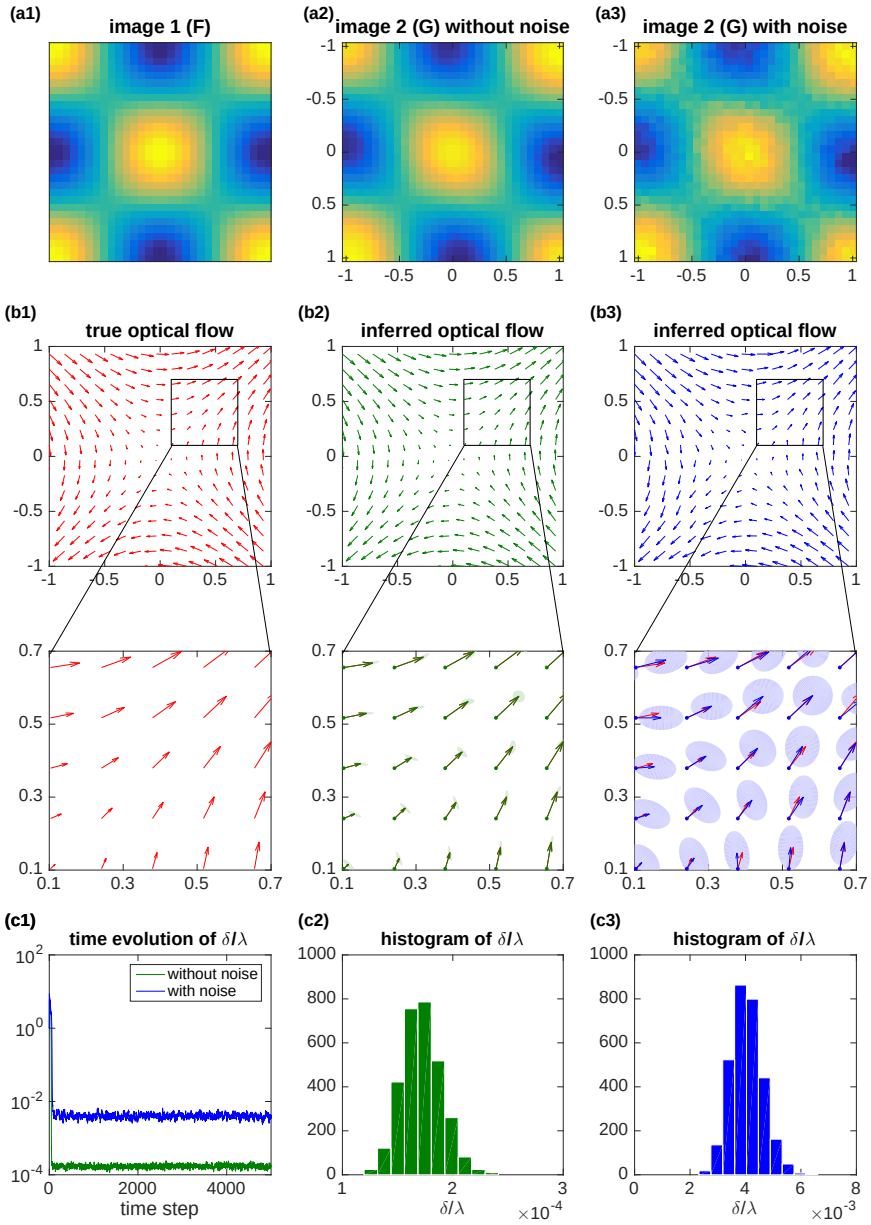

FIGURE 3. Statistical inversion based optical flow example 1: $\langle U(x,y), V(x,y)\rangle = \langle x, y\rangle$. Top row (a1-a3): image data generated according to Eq. (37) for F (a1) and using Eq. (38) for G either without noise (a2) or with noise (a3). Middle rows (b1-b3): the true optical flow field (b1) and the inferred optical flow field from the MCMC samples. In all panels corresponding to the inferred optical flow, the mean vector at each point is plotted as an arrow; for the zoomed-in views, we show the computed confidence region defined by Eq. (39) as shaded ellipses around the tip of the arrows. Bottom row (c1-c3): time evolution (c1) as well as the distribution (c2-c3) of the effective regularization parameter σ/λ, where the distributions are obtained after discarding the initial transient in the MCMC sampling process.

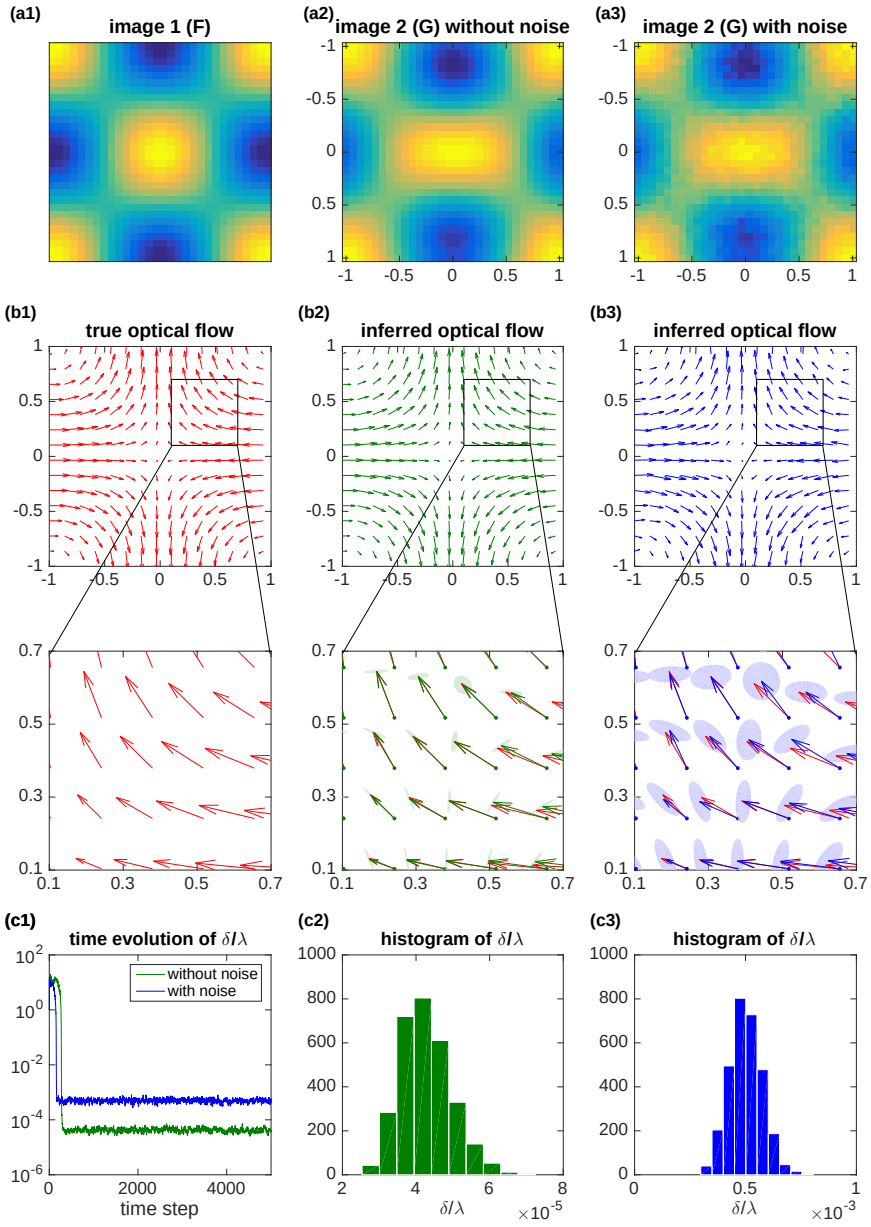

FIGURE 4. Statistical inversion based optical flow example 1: $\langle U(x,y), V(x,y)\rangle = \langle x, y\rangle$. Top row (a1-a3): image data generated according to Eq. (37) for F (a1) and using Eq. (38) for G either without noise (a2) or with noise (a3). Middle rows (b1-b3): the true optical flow field (b1) and the inferred optical flow field from the MCMC samples. In all panels corresponding to the inferred optical flow, the mean vector at each point is plotted as an arrow; for the zoomed-in views, we show the computed confidence region defined by Eq. (39) as shaded ellipses around the tip of the arrows. Bottom row (c1-c3): time evolution (c1) as well as the distribution (c2-c3) of the effective regularization parameter σ/λ, where the distributions are obtained after discarding the initial transient in the MCMC sampling process.

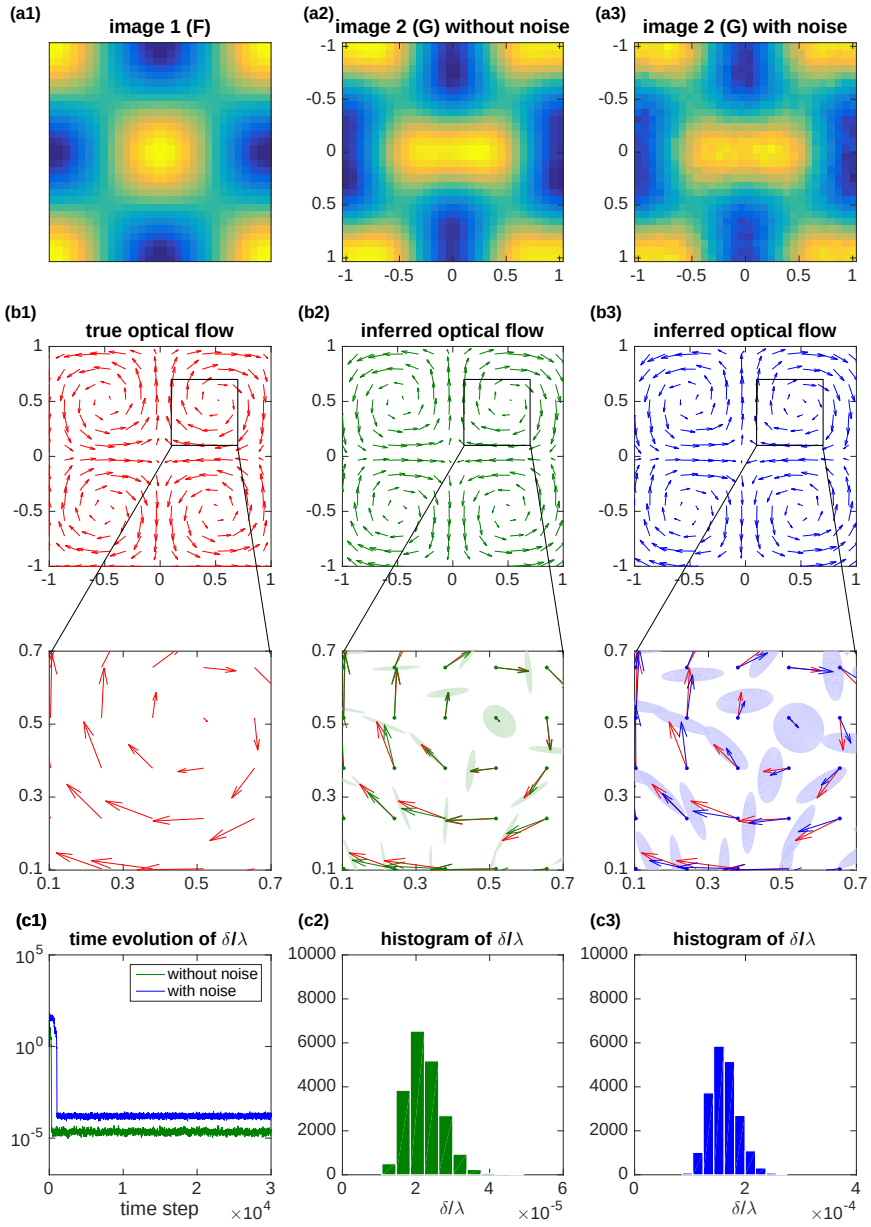

FIGURE 5. Statistical inversion based optical flow example 1: $\langle U(x,y), V(x,y)\rangle = \langle x, y\rangle$. Top row (a1-a3): image data generated according to Eq. (37) for F (a1) and using Eq. (38) for G either without noise (a2) or with noise (a3). Middle rows (b1-b3): the true optical flow field (b1) and the inferred optical flow field from the MCMC samples. In all panels corresponding to the inferred optical flow, the mean vector at each point is plotted as an arrow; for the zoomed-in views, we show the computed confidence region defined by Eq. (39) as shaded ellipses around the tip of the arrows. Bottom row (c1-c3): time evolution (c1) as well as the distribution (c2-c3) of the effective regularization parameter σ/λ, where the distributions are obtained after discarding the initial transient in the MCMC sampling process.